Preserving Patient Privacy in Dynamic Treatment Regimes

Dylan Spicker Erica E.M. Moodie Susan M. Shortreed

Department of Epidemiology, Biostatistics and Occupational Health
McGill University

Tuesday May 30, 2023

Precision Medicine

Treat the patient, not the disease.

Individualized Treatment Rules: Mathematical Statement

We want to estimate a decision function,

$$d\colon \mathcal{H} \longrightarrow \mathcal{A} = \{0, 1\},$$

where $H \in \mathcal{H}$ is the patient history and $A \in \mathcal{A}$ is the treatment decision.

We call this function a individualized treatment rule (ITR).

ITR Estimation as Classification

An ITR, d, has value

$$V(d) = E\{E[R|A = d(H)]\}.$$

Optimal ITRs maximize the value.

ITR Estimation as Classification

An ITR, d, has value

$$V(d) = E\{E[R|A = d(H)]\}.$$

Optimal ITRs maximize the value. Optimal ITRs minimize

$$E[R|A = 1] + E[R|A = -1] - V(d) = E\left[\frac{R}{P(A|H)} I(A \neq d(H))\right].$$

Outcome-Weighted Learning (OWL) estimates optimal ITRs by minimizing a regularized, empirical version of this error.

Support Vector Machines (SVM)

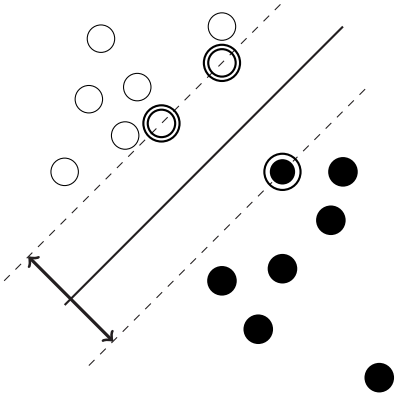

SVMs use hyperplanes to solve classification problems.

The resulting classifier exists as

$$f(H) = \sum_{i \in \mathcal{SV}} \alpha_i A_i K(H_i, H).$$

Privacy Implications of SVMs

Generally, the resulting decision function requires the **direct release** of the support vectors.

$$f(H) = \sum_{i \in \mathcal{SV}} \alpha_i A_i \exp\left(-\sigma^2 \|H_i - H\|\right)$$

Differential Privacy

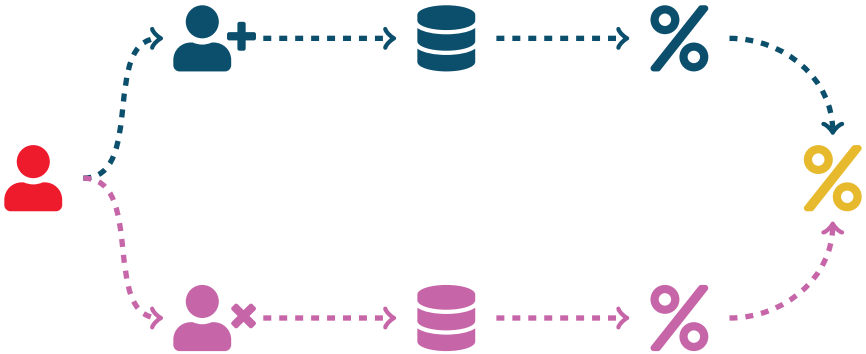

Differential Privacy

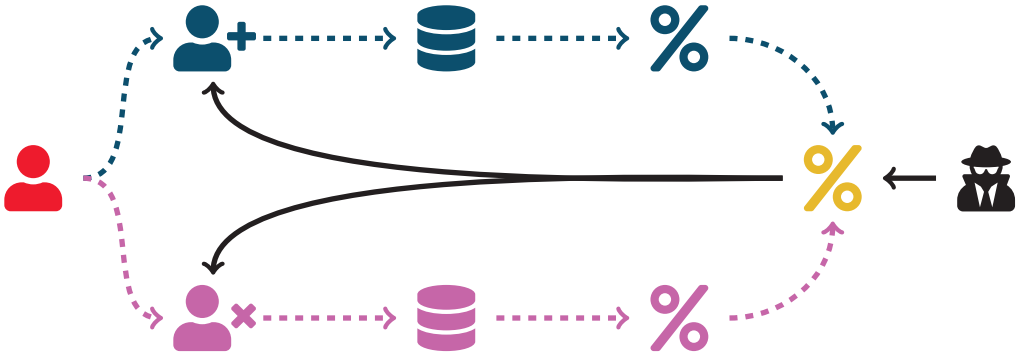

Differential Privacy: Mathematical Statement

We say that an estimator, $\mathcal{M}$, is ϵ-differentially private if for **all** neighbouring datasets, $\mathbb{X}$ and $\mathbb{X}^\dagger$, we have:

$$\frac{P(\mathcal{M}(\mathbb{X}) \in \mathcal{Y})}{P(\mathcal{M}(\mathbb{X}^\dagger) \in \mathcal{Y})} \leq e^{\epsilon}.$$

Private Outcome-Weighted Learning (PrOWL)

We propose a differentially private implementation of OWL, called PrOWL.

1. Approximate the kernel in finite dimensions.

Private Outcome-Weighted Learning (PrOWL)

We propose a differentially private implementation of OWL, called PrOWL.

1. Approximate the kernel in finite dimensions.
2. Compute the standard OWL estimator.

Private Outcome-Weighted Learning (PrOWL)

We propose a differentially private implementation of OWL, called PrOWL.

1. Approximate the kernel in finite dimensions.
2. Compute the standard OWL estimator.
3. Perturb the vector with Laplace distributed errors.

PrOWL Guarantees

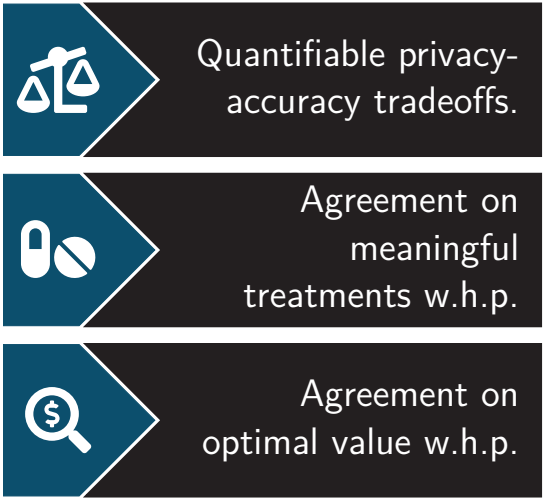

Summary

Privacy should be a major concern within precision medicine and beyond.

Differential privacy provides one framework for addressing these concerns, with promising results thus far.

Thank You!

www.dylanspicker.com | dylan.spicker@mcgill.ca